arXiv:0711.0458v1 [stat.ME] 3 Nov 2007

Bayesian finite mixtures: a note on prior specification and posterior computation

Agostino Nobile*
University of Glasgow, Scotland.

May 2005

Abstract

A new method for the computation of the posterior distribution of the number k of components in a finite mixture is presented. Two aspects of prior specification are also studied: an argument is made for the use of a $Poi(1)$ distribution as the prior for k; and methods are given for the selection of hyperparameter values in the mixture of normals model, with natural conjugate priors on the components parameters.

Keywords: Galaxy data, Marginal Likelihood, Markov Chain Monte Carlo, Mixtures of Normals.

1 Introduction

Finite mixture distributions have become widely used as a tool of semi-parametric inference: they partake of the conceptual simplicity of parametric models and of the flexibility of non-parametric ones. This paper is a contribution to the Bayesian analysis of finite mixtures with an unspecified number of components. I give arguments to support the use of a $Poi(1)$ prior for the number of components and present a new method for the numerical computation of its posterior. The method exploits a fundamental probability identity already used by Chib (1995), but combines it with the representation of mixture marginal likelihoods given in Nobile (2004). I also discuss a more specific topic, hyperparameter selection in a finite mixture of univariate normals. Throughout the paper, the galaxy data set is used for illustrative purposes.

The remainder of this section provides a brief introduction to Bayesian finite mixtures, representations of the associated marginal likelihoods and the mixture of normals model. Section 2 deals with the estimation of marginal likelihoods using the frequency of empty components in a Markov Chain Monte Carlo sample of the mixture allocations. Section 3 argues that the structure of the model suggests the $Poi(1)$ distribution as a suitable prior for the number of components, when no substantive information on it is available. Section 4 concerns the more practical issue of hyperparameter determination in the mixture of normals model, when natural conjugate priors on the means and variances of the components are employed.

*Department of Statistics, University of Glasgow, Glasgow G12 8QW, Scotland.
Email: agostino@stats.gla.ac.uk

1.1 Bayesian finite mixtures

A finite mixture is a distribution with density, with respect to some underlying measure, given by

$$f(x) = \sum_{j=1}^{k} \lambda_j p_j(x|\theta_j). \tag{1}$$

The weights λ_j are non-negative and sum to 1, while the component densities $p_j(\cdot|\theta_j)$ belong to some known parametric family. Observations $x_1, \ldots, x_n$ are regarded as proceeding from the distribution (1) and interest lies in the number of components k and, conditional on k, in the weights λ and the components' parameters θ.

The model can be rewritten by introducing latent allocation vectors $g = (g_1, \ldots, g_n)$ with $g_i \in \{1, \ldots, k\}$ denoting the mixture component that generated the i-th observation:

$$\begin{array}{rcll} \Pr[g_i = j|k, \lambda] & = & \lambda_j & i = 1, \ldots, n, \quad \text{independently} \\ x_i|k, g, \lambda, \theta & \overset{\text{ind.}}{\sim} & p_j(x_i|\theta_j), & j = g_i, \quad i = 1, \ldots, n. \end{array} \tag{2}$$

In the Bayesian analysis of the model, typically one assumes that $\lambda|k \sim Dir(\alpha_1, \ldots, \alpha_k)$, where the α_j's are fixed constants. Also, the component parameters θ_j are assumed a priori independent, conditionally on k and, possibly, a vector of hyperparameters ϕ:

$$\pi(\theta|k, \phi) = \prod_{j=1}^{k} \pi_j(\theta_j|\phi_j).$$

If a prior distribution $\pi(k)$ is specified, then one can obtain a sample from the joint posterior of (k, λ, θ) by means of Markov chain Monte Carlo methods, see e.g. Richardson and Green (1997), Phillips and Smith (1996), Stephens (2000a), Nobile and Fearnside (2005). Inference about λ and θ is not straightforward, because the likelihood is invariant with respect to permutations of the components' labels. Achieving identifiability by imposing constraints on the parameters does not always work and other methods have been proposed, see Richardson and Green (1997) and its discussion (especially the contributions of G. Celeux and M. Stephens), Celeux, Hurn and Robert (2000), Stephens (2000b), Frühwirth-Schnatter (2001), Nobile and Fearnside (2005).

An alternative to sampling from the posterior of (k, λ, θ) consists of estimating the marginal likelihoods f_k of the mixture model with k components:

$$f_k := f(x|k) = \iint f(x|k, \lambda, \theta)\pi(\lambda, \theta|k)\, \mathrm{d}\lambda\, \mathrm{d}\theta, \qquad k = 1, 2, \ldots, k_{\max}.$$

Each marginal likelihood estimate makes use of MCMC output for a model with fixed k. The estimates can be used to compute Bayes factors for k vs. $k-1$ components, or to compute the posterior of k, $\pi(k|x) \propto \pi(k)f_k$. It should be noted that estimation of the marginal likelihood from MCMC output is not as simple as other posterior inference using MCMC, and as a consequence several methods have been proposed, see e.g. Chib (1995), Raftery (1996), DiCiccio et al. (1997), Gelman and Meng (1998) and references therein.

1.2 Marginal likelihoods of finite mixtures

The marginal likelihoods can be rewritten as

$$f_k = \sum_{g \in \mathcal{G}_k} f(g|k) f(x|k, g) \tag{3}$$

where the sum extends over the set $\mathcal{G}_k$ of all the allocation vectors with entries less than or equal to k, see e.g. Nobile (1994, 2004). In equation (3), $f(g|k) = \int f(g|k,\lambda)\pi(\lambda|k)\,\mathrm{d}\lambda = \frac{\Gamma(\alpha_{0k})}{\Gamma(\alpha_{0k}+n)}\prod_{j=1}^{k}\frac{\Gamma(\alpha_j+n_j)}{\Gamma(\alpha_j)}$ with $\alpha_{0k} = \sum_{j=1}^{k}\alpha_j$ and n_j equal to the number of observations that g allocates to component j. The other term in the right hand side of (3), $f(x|k,g)$, is obtained by integrating $f(x|k,g,\theta)$ from (2) with respect to the prior distribution of $\theta|k$. Although this integration can be performed in closed form only for some prior distributions, notably natural conjugate priors on θ, representation (3) is always valid. Under the assumption that the Dirichlet hyperparameters α_j and the prior distributions $\pi_j(\cdot|\phi_j)$ remain the same for fixed j as k varies, the marginal likelihoods enjoy further representations. Partition the set of allocation vectors $\mathcal{G}_k$ as

$$\mathcal{G}_k = \bigcup_{t=1}^{k}\mathcal{G}_t^\star, \qquad \mathcal{G}_t^\star \cap \mathcal{G}_s^\star = \emptyset, \qquad t \neq s$$

where $\mathcal{G}_t^\star$ is the set of allocation vectors which assign at least one observation to component t and none to higher components. Also, let $f_t^\star$ be the portion of the marginal likelihood f_t that accounts for vectors g allocating at least one observation to component t and none to higher:

$$f_t^\star = \sum_{g\in\mathcal{G}_t^\star} f(g|t)f(x|t,g). \tag{4}$$

Then one can show (see Nobile 2004, page 2049) that, for all $g \in \mathcal{G}_t^\star$ with $t < k$,

$$f(x|k,g) = f(x|t,g) \qquad \text{and} \qquad \frac{f(g|k)}{f(g|t)} = \frac{\Gamma(\alpha_{0k})}{\Gamma(\alpha_{0k}+n)}\frac{\Gamma(\alpha_{0t}+n)}{\Gamma(\alpha_{0t})} =: a_{kt} \tag{5}$$

and that

$$f_k = \sum_{t=1}^{k} a_{kt}f_t^\star \tag{6}$$

$$= a_{k,k-1}f_{k-1} + f_k^\star. \tag{7}$$

If the prior distribution of $(\lambda,\theta)|k$ is invariant to permutations of the components labels, a stronger result is available. Let $\mathcal{G}_h^t$ be the subset of $\mathcal{G}_t^\star$ consisting of allocations with $h \leq t$ non-empty components. In particular, any vector $g \in \mathcal{G}_h^h \subset \mathcal{G}_h$ assigns at least one observation to each mixture component $1,\ldots,h$. Let $f_h^\dagger$ be the portion of f_h which corresponds to allocations with no empty components:

$$f_h^\dagger = \sum_{g\in\mathcal{G}_h^h} f(g|h)f(x|h,g).$$

Then (Nobile 2004, page 2053)

$$f_k = \sum_{h=1}^{k\wedge n}\binom{k}{h}a_{kh}f_h^\dagger. \tag{8}$$

For related representations see Ishwaran, James and Sun (2001).

1.3 Mixtures of univariate normals

The method to be presented in the following section is of general applicability. Since mixtures of univariate normals will be used as an illustration, I introduce here some notation. It is assumed that the component densities $p_j(\cdot|\theta_j)$ are normal with mean m_j and variance r_j^{-1}:

$$x_i|k,g,\lambda,\theta \overset{ind.}{\sim} N(m_j, r_j^{-1}), \qquad j = g_i, \quad i = 1,\ldots,n.$$

Independent natural conjugate priors are placed on $\theta_j = (m_j, r_j)$, $j = 1, \ldots, k$:

$$\begin{aligned} r_j &\overset{\text{ind.}}{\sim} Ga(\gamma, \delta) \\ m_j | r_j &\overset{\text{ind.}}{\sim} N(\mu, \{\tau r_j\}^{-1}), \end{aligned} \tag{9}$$

with $E(r_j) = \gamma/\delta$. See Diebolt and Robert (1994) or Nobile and Fearnside (2005) for more details. Other priors on θ, such as the independent prior used by Richardson and Green (1997), could be used as well.

The prior distribution (9) requires the specification of four hyperparameters. The overall mean μ is set to a round value close to the sample mean $\overline{x}$. The shape parameter γ is half the degrees of freedom of the prior predictive t distribution. I choose $\gamma = 2$, to have a t_4 prior predictive, with relatively thick tails, but finite second order moments. The choice of the scale parameter δ and of τ, the prior ratio between within components variance and variance of the means, is discussed in Section 4.

2 Marginal likelihoods from empty components

For a given parametric model $f(x|\theta)$ and prior distribution $f(\theta)$, the marginal likelihood of the observed data x is defined as $f(x) = \int f(x|\theta) f(\theta) \, \mathrm{d}\theta$. Using Bayes theorem, $f(x)$ can be rewritten as

$$f(x) = \frac{f(\theta) f(x|\theta)}{f(\theta|x)} \tag{10}$$

where $f(\theta)$ and $f(x|\theta)$ are assumed computable, including their normalizing constants, and the formula holds for any parameter value θ. Expression (10) forms the basis of a method of marginal likelihood estimation, see Chib (1995) and Raftery (1996). In short, although typically the posterior $f(\theta|x)$ cannot be evaluated exactly, an estimate of it at some parameter value θ can be obtained using a Monte Carlo sample; substituting this estimate into (10) yields an estimate of $f(x)$.

In the context of Section 1.2, the marginal likelihood for the model with k components can be written as

$$f_k = \frac{f(g|k) f(x|k, g)}{f(g|k, x)}. \tag{11}$$

Here the allocation vector g plays the role of θ in the above discussion and everything is conditional on k. Since (11) holds for all $g \in \mathcal{G}_k$, it still holds if one sums both numerator and denominator over any non-empty set E:

$$f_k = \frac{\sum_{g \in E} f(g|k) f(x|k, g)}{\sum_{g \in E} f(g|k, x)}. \tag{12}$$

Letting $E = \mathcal{G}_k^\star$, the denominator of (12) is the posterior probability of $\mathcal{G}_k^\star$, while the numerator equals $f_k^\star = f_k - a_{k,k-1} f_{k-1}$, using equations (4) and (7). One then obtains

$$f_k = \frac{f_k - a_{k,k-1} f_{k-1}}{\Pr[\mathcal{G}_k^\star | k, x]}$$

and after rearranging

$$\frac{f_k}{f_{k-1}} = \frac{a_{k,k-1}}{1 - \Pr[\mathcal{G}_k^\star | k, x]}. \tag{13}$$

The left hand side of (13) is the Bayes factor $B_{k,k-1}$ for the model with k components against the model with $k-1$ components. In the right hand side $a_{k,k-1}$ is a known constant, while the denominator is the posterior probability, according to the model with k components, that the k-th component is empty, which can be easily estimated using a MCMC sample from $f(g|k, x)$. In

some mixture models $f_1 = f_1^\star$ is computable exactly; if this is the case, estimates of the marginal likelihoods, if needed, can be readily produced from the Bayes factors. Otherwise, one can still obtain estimates of normalized marginal likelihoods, by setting $f_1 = 1$ and then rescaling the sequence of f_k's.

Using formula (13) is somewhat wasteful, since it only employs the fixed k MCMC sample to estimate $\Pr[\mathcal{G}_k^\star|k,x]$: the MCMC sample for $k+1$ components can be used to estimate $\Pr[\mathcal{G}_k^\star|k+1,x]$, a quantity that is related to the probability in (13). Let $\Pr[\mathcal{G}_t^\star|k,x]$ with $t \leq k$ be the posterior probability, conditional on k components, that component t is non-empty and components $t+1$ through k are empty. One can show, see the Appendix, that

$$\frac{f_{t+1}^\star}{f_t^\star} = a_{t+1,t}\, \frac{\sum_{k=t+1}^{k_{\max}} Pr[\mathcal{G}_{t+1}^\star|k,x]}{\sum_{k=t+1}^{k_{\max}} Pr[\mathcal{G}_t^\star|k,x]}. \tag{14}$$

Setting $f_1^\star = f_1$ if available, or $f_1 = 1$ if not, the sequence $f_t^\star$ $t = 1,\ldots,k_{\max}$ can be estimated by replacing the probabilities in (14) with MCMC estimates. An application of (6), followed by rescaling, then produces estimates of the normalized marginal likelihoods.

Formulae (13) and (14) do not assume that the prior on $(\lambda,\theta)|k$ is invariant to permutations of the components labels, only that the hyperparameters α_j, ϕ_j are the same for all $k \geq j$. If the prior is invariant, the additional symmetry can be exploited as follows. Let $\widetilde{\mathcal{G}}_h^k = \bigcup_{t=h}^k \mathcal{G}_h^t$ be the set of allocations g in $\mathcal{G}_k$ which assign observations to exactly h components. Then $\Pr[\widetilde{\mathcal{G}}_h^k|k,x]$ is the posterior probability, conditional on k components, that $h \leq k$ components are non-empty. One can show, see the Appendix, that

$$\frac{f_{h+1}^\dagger}{f_h^\dagger} = (h+1)\, a_{h+1,h}\, \frac{\sum_{k=h+1}^{k_{\max}} \Pr[\widetilde{\mathcal{G}}_{h+1}^k|k,x]}{\sum_{k=h+1}^{k_{\max}} (k-h)\Pr[\widetilde{\mathcal{G}}_h^k|k,x]}. \tag{15}$$

Replacing the probabilities in (15) with MCMC estimates and setting $f_1^\dagger$ to 1 (or f_1 if available), yields estimates of the sequence of $f_h^\dagger$'s; plugging these estimates in formula (8) and rescaling produces estimates of normalized marginal likelihoods.

To illustrate the method, formula (15) was used to compute the marginal likelihood of k components for the galaxy data. This data set consists of velocity measurements (1000 Km/sec) of 82 galaxies from the Corona Borealis region. Since its appearance in Roeder (1990), it has been studied by several authors, see Aitkin (2001) for an interesting comparison of likelihood and Bayesian analyses of this data set. The data was modelled as a finite mixture of univariate normals, as set out in Section 1.3. The weights hyperparameters α were set to 1, while the other hyperparameters were $\mu = 20$, $\tau = 0.04$, $\gamma = 2$ and $\delta = 2$, their choice is discussed in Section 4. In this example I used Gibbs sampling of the allocation vectors g, after integrating out the weights and components parameters, see Nobile and Fearnside (2005). However, the method applies equally well to the Gibbs sampling scheme involving both parameters and allocations, see for instance Diebolt and Robert (1994) and Richardson and Green (1997), as long as empty components are allowed. Each Gibbs sampler with fixed k was run for 20000 sweeps, with 1000 sweeps of burn-in. The final allocation in the run with k components served as the starting allocation for the run with $k+1$ components. The estimates of the marginal likelihoods normalized to sum to 1 are displayed as

line-joined dots in Figure 1. For comparison, the figure also contains the estimate of the posterior of k with uniform prior on $k = 1, \ldots, k_{\max} = 50$ using a different method, the allocation sampler of Nobile and Fearnside (2005). This sampler was run for 1 million sweeps with a burn-in of 10000 sweeps and keeping only one draw every 10. The agreement between the estimates from the two

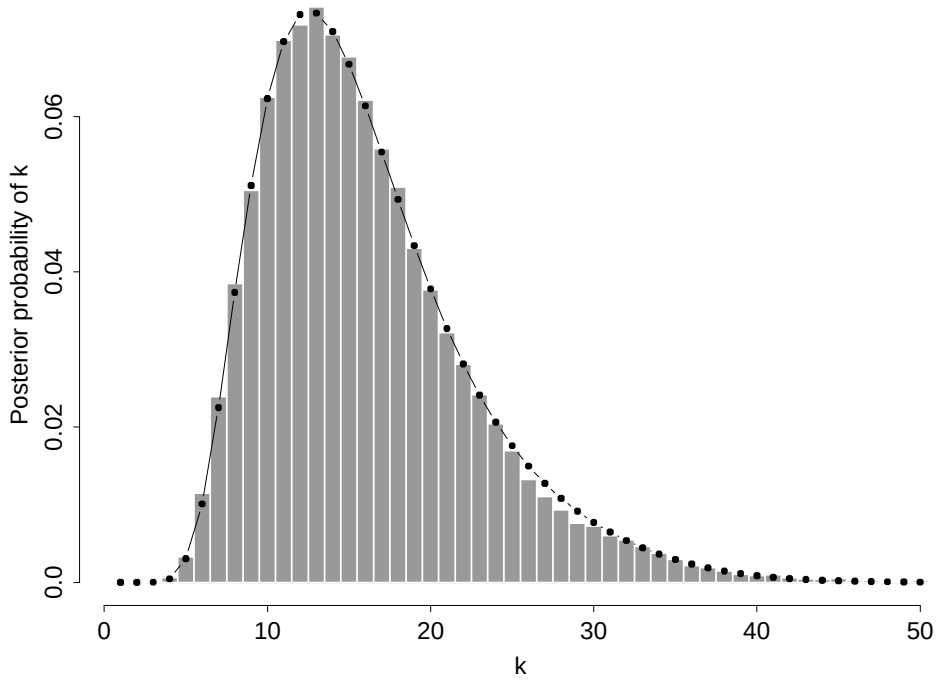


FIGURE 1: Two estimates of the posterior distribution of k for the galaxy data, using a discrete uniform prior on $k = 1, \ldots, k_{\max} = 50$. Histogram gives the frequency of model with k components using the allocation sampler of Nobile and Fearnside (2005). Line-joined dots are the normalized marginal likelihoods using formula (15). See main text for hyperparameter values used.

unrelated methods provides a welcome check on them, all the more so since visual inspection of the galaxy data suggests between three and six clusters, while the posterior of k displayed in Figure 1 assigns to this range of values a probability smaller than 0.02. If one is to believe the estimates in Figure 1, as the agreement between the two methods seems to suggest, it would seem that the posterior of k has little to tell about the number of clusters in a data set. In the next section I argue that this is not the case and that replacing the uniform prior on k with a $Poi(1)$ distribution yields a posterior that is more suitable for inference about the number of actual groups in the data.

3 The prior distribution of the number of components

In this section I assume that the prior on $(\lambda, \theta)|k$ is invariant to permutations of the components labels. Recall from Section 1.2 that $f_h^\dagger$ is the part of the marginal likelihood f_h corresponding to no empty components

$$f_h^\dagger = \sum_{g \in \mathcal{G}_h^h} f(g|h) f(x|h, g)$$

and that representation (8) holds:

$$f_k = \sum_{h=1}^{k \wedge n} \binom{k}{h} a_{kh} f_h^\dagger.$$

To have an understanding of how formula (8) arises, look at Figure 2 which displays the nested structure of $\mathcal{G}_4$, with each set of digits denoting allocation vectors with entries equal to those digits only. From formula (3), f_4 is the sum over $\mathcal{G}_4$ of $f(g|k=4)f(x|k=4,g)$. Formula (8) gives this

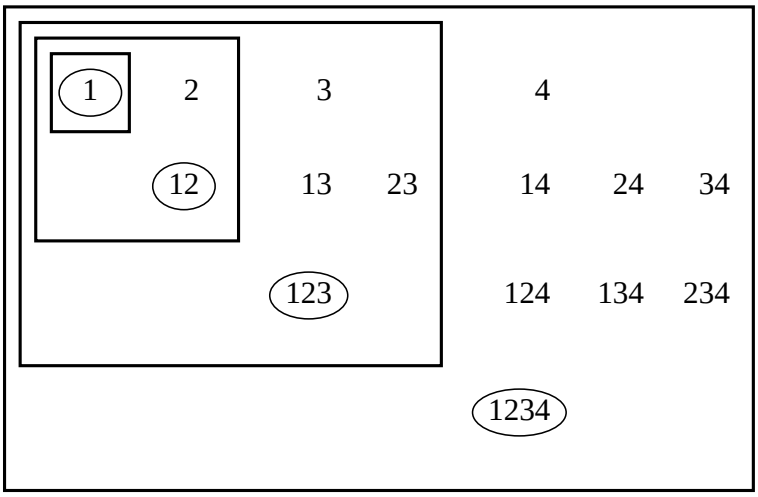


FIGURE 2: The nested structure of $\mathcal{G}_k$, $k=4$. Each set of digits denotes a set of allocation vectors with entries equal to those digits only. Nested boxes denote $\mathcal{G}_1, \ldots, \mathcal{G}_4$. Ellipses enclose the sets $\mathcal{G}_h^h$, $h=1,\ldots,4$. In box $\mathcal{G}_k$, the complement of the enclosed $\mathcal{G}_{k-1}$ box is the set $\mathcal{G}_k^\star$.

sum in terms of $f_h^\dagger$'s, which are sums over the sets $\mathcal{G}_h^h$, $h=1,\ldots,4$ denoted by ellipses in Figure 2. The terms a_{kh} serve to rescale $f(g|h)$ to $f(g|k)$, while the combinatorial terms $\binom{k}{h}$ give the number of "copies" of $\mathcal{G}_h^h$ that are present in $\mathcal{G}_k$.

A consequence of formula (8) is that the marginal likelihood of k components may derive to a large extent from allocations with less than k non-empty components. For instance, consider a hypothetical data set of $n=80$ observations clearly clustered in nine well separated groups, to such an extent that $f_h^\dagger / f_9^\dagger$ is nearly 0, for $h \neq 9$. Then formula (8) implies that

$$\frac{f_k}{f_9} = \binom{k}{9} a_{k9}, \qquad k \geq 9. \tag{16}$$

With $\alpha = 1$, one obtains the values reported in Table 1. With a discrete uniform prior on k,

k	9	10	11	12	13	14	15
f_k/f_9	1.	1.011	0.618	0.299	0.127	0.050	0.018

TABLE 1: Ratio of marginal likelihood f_k/f_9 for a hypothetical data set with $n=80$ observations, such that $f_h^\dagger / f_9^\dagger = 0$, $h \neq 9$.

the posterior of k gives probability less that 1/3 to $k=9$. Put differently, upper bounds on the posterior of k can be derived from representation (8). Table 2, taken from Nobile (2004), displays upper bounds corresponding to a discrete uniform prior and $\alpha = 1$. Nobile (2004) contains further discussion and tables for $\alpha = 2$ and $\alpha = 0.5$. The overall conclusion is that the bounds are weaker for larger sample sizes, smaller values of k and larger values of α.

n	k									
	1	2	3	4	5	6	7	8	9	10
20	0.9000	0.7286	0.5299	0.3456	0.2880	0.2419	0.1954	0.1756	0.1505	0.1335
50	0.9600	0.8847	0.7826	0.6645	0.5414	0.4233	0.3175	0.3119	0.2835	0.2402
100	0.9800	0.9412	0.8858	0.8170	0.7385	0.6541	0.5677	0.4828	0.4023	0.3322
500	0.9960	0.9880	0.9762	0.9607	0.9417	0.9193	0.8938	0.8656	0.8350	0.8022

TABLE 2: Bounds on $\pi(k|x)$ for several sample sizes n, $\pi(k) = 1/k_{max}, k = 1, \ldots, k_{max} = 50$, $\alpha = 1$.

It is worth mentioning at this point that, as the sample size grows, the marginal likelihood of k components will tend to reflect more and more only allocations with no empty components. Formally,

$$\binom{k}{h} a_{kh} \to \delta_{\{k=h\}}, \qquad n \to \infty, \tag{17}$$

see the Appendix for a proof. Hence, from formula (8), $f_k - f_k^\dagger \to 0$ as $n \to \infty$.

Returning to the example of nine well separated groups, it is the combinatorial term $\binom{k}{h}$ that makes $f_{10} > f_9$ in Table 1. If one were to drop the $\binom{k}{9}$ term from equation (16), the entries in Table 1 would be as in Table 3. The $\binom{k}{9}$ term accounts for the fact that with $k = 10$ components,

k	9	10	11	12	13	14	15
f_k/f_9	1.	0.10112	0.01124	0.00136	0.00018	0.00002	0.00000

TABLE 3: Ratios f_k/f_9 as in Table 1, but dropping the term $\binom{k}{9}$ from equation (16).

there are ten possible ways of choosing nine components to have observations and one component to be empty. Of course, this is a consequence of the model entertained and its ability to allow for empty components, which correspond to mass on small values for some weights in the prior of λ. Nonetheless, the increasing effect on the marginal likelihoods, as k grows, of the many ways in which some of k components may be empty, is a rather unappealing feature of the model. Nobile (2004) has suggested to shift attention from the number of components to the number of non-empty components and to compute its posterior distribution. In this paper I follow a different approach: trying to counteract the combinatorial terms in the marginal likelihoods by an appropriate choice of the prior distribution of k.

Multiplying equation (8) by the prior distribution $\pi(k)$ and writing the result explicitly for the first few k, one has

$$\begin{aligned}
\pi(1|x) &= A \cdot \pi(1) f_1^\dagger \\
\pi(2|x) &= A \left[\pi(2) f_2^\dagger + \pi(2) \binom{2}{1} a_{21} f_1^\dagger \right] \\
\pi(3|x) &= A \left[\pi(3) f_3^\dagger + \pi(3) \binom{3}{2} a_{32} f_2^\dagger + \pi(3) \binom{3}{1} a_{31} f_1^\dagger \right] \\
&\ldots
\end{aligned}$$

where A is a normalizing constant. Although there is no prior $\pi(k)$ which exactly cancels out the binomial coefficients $\binom{k}{h}$, one can keep the contribution of $f_h^\dagger$ to $\pi(k|x)$ small, relative to its

contribution to $\pi(h|x)$, by requiring that

$$\pi(h) = \sup_{k>h} \left\{ \pi(k) \binom{k}{h} \right\} \qquad h = 1, 2, \ldots \ . \tag{18}$$

It is easy to verify that a $Poi(1)$ distribution satisfies equations (18). Indeed, every prior $\pi(k)$ satisfying equations (18) is proportional to a truncated $Poi(1)$ distribution, see the Appendix. For simplicity, I will take the prior on k to be $Poi(1)$.

One way to illustrate the effect of the $Poi(1)$ prior on k is by recomputing the bounds on $\pi(k|x)$ with this prior; they are reported in Table 4. Compared to the bounds with a discrete uniform prior in Table 2, they are much weaker, especially for higher values of k.

n	k									
	1	2	3	4	5	6	7	8	9	10
20	0.9525	0.9114	0.8756	0.8441	0.8162	0.7913	0.7690	0.7488	0.7306	0.7140
50	0.9804	0.9619	0.9445	0.9280	0.9124	0.8976	0.8836	0.8703	0.8576	0.8455
100	0.9901	0.9805	0.9712	0.9621	0.9533	0.9447	0.9364	0.9283	0.9204	0.9128
500	0.9980	0.9960	0.9940	0.9921	0.9901	0.9882	0.9863	0.9844	0.9825	0.9806

TABLE 4: Bounds on $\pi(k|x)$ for several sample sizes n, $\pi(k) \propto 1/k!$, $\alpha = 1$.

For another illustration, reconsider the example of nine well separated groups in Table 1. The ratio of posterior probabilities $\pi(k|x)/\pi(9|x)$ using a $Poi(1)$ prior are given in Table 5.

k	9	10	11	12	13	14	15
$\pi(k\|x)/\pi(9\|x)$	1.	0.10112	0.00562	0.00023	0.00001	0.00000	0.00000

TABLE 5: Ratio of posterior probabilities $\pi(k|x)/\pi(9|x)$ for the same hypothetical data set as in Table 1, using a $Poi(1)$ prior on k.

As a further illustration, return to the galaxy data example in Section 2. Figure 3 contains the posterior of k computed using the same hyperparameters and methods as in Figure 1, but with a $Poi(1)$ prior on k, rather than discrete uniform. Other examples, for real and artificial data sets, of posterior distributions of k based on a $Poi(1)$ prior can be found in Nobile and Fearnside (2005).

4 Mixtures of normals: hyperparameter selection

This section is concerned with the choice of hyperparameters in mixtures of univarite normals, with natural conjugate priors on the means and variances. The method to be described can be readily adapted to the case of multivariate components, or components from other parametric families. I continue to use the galaxy data set for illustrative purposes. The marked sensitivity to prior specification exhibited in the analysis of this data is, in my experience, far from typical. However, it demonstrates well what difficulties may arise. Patterns of dependence of the marginal likelihood on the prior of θ are likely to be simpler in one-parameter families; see Aitkin (2001, page 289) for a related remark.

Figure 4 displays estimates of the posterior distribution of the number of components for the galaxy data, corresponding to several values of the hyperparameters τ and δ. The other hyperparameters were set to $\mu = 20$ and $\gamma = 2$, as discussed in Section 1.3. The prior on k was $Poi(1)$ and

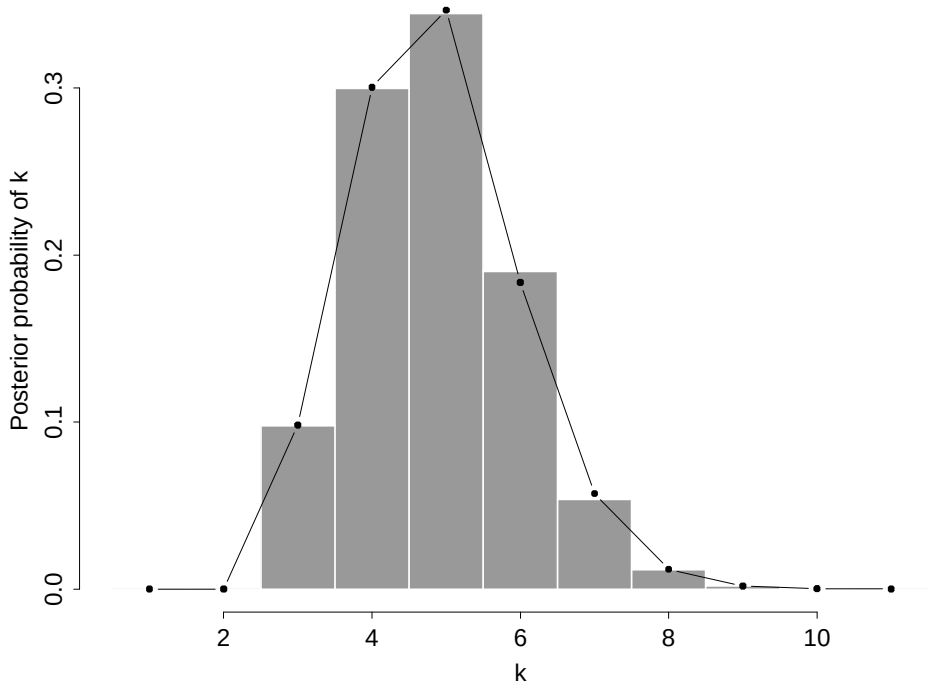


FIGURE 3: Two estimates of the posterior distribution of k for the galaxy data, using a Poisson(1) prior. Histogram gives the frequency of model with k components using the allocation sampler of Nobile and Fearnside (2005). Solid line is the posterior of k obtained from the normalized marginal likelihood using formula (15). See main text for hyperparameter values used.

computations were done using formula (15). Although in all plots most of the mass is concentrated on values of k between 2 and 8, a simple glance at the figure conveys how dependent on hyperparameter values $\pi(k|x)$ may be. One can also see that, for given δ, as τ increases at first posterior mass shifts to higher values of k, and then it moves back to lower values of k. The behaviour for τ fixed and δ increasing consists, apart for few exceptions, of a shift of probability mass from higher to lower values of k. Most pairs (τ, δ) yield negligible posterior mass for $k = 2$. However, some pairs in the upper right corner of the plot assign considerable mass to it. These pairs correspond to a prior distribution that makes likely high values of the variance within each normal component; in turns this makes it plausible to place in a single group the smallest and largest observations in the galaxy data, with a central group accounting for most of the other observations.

Putting a hyperprior $\pi(\tau, \delta)$ on the two hyperparameters, and sampling from the joint posterior of all the unknowns, including τ and δ, did not solve the problem. Some experimentation with a few hyperpriors showed that $\pi(k|x)$ was to a considerable extent affected by the choice of hyperprior: the marginal posterior distributions of τ and δ had very long tails and changed markedly with $\pi(\tau, \delta)$. For this reason, I preferred to adopt an empirical Bayes stance and estimate τ and δ rather than mixing with respect to their posterior distribution.

I settled on independent priors: $Un(0, 1)$ for $(1 + \tau)^{-1}$, the prior proportion of variance within a component to the total variance, and $Un(0, \delta_U)$ for δ, where $\delta_U = (\gamma - 1)s_x^2$ and s_x^2 is the sample variance. The choice of δ_U yields a prior expectation of the components variance equal to $s_x^2/2$. The random walk Metropolis-Hastings algorithm was used to update τ and δ given all other variables. Figure 5 displays boxplots of the marginal posterior distributions of τ and δ, on a logarithmic scale, conditional on k. Both plots display a pattern whereby a clear change of level occurs as k increases.

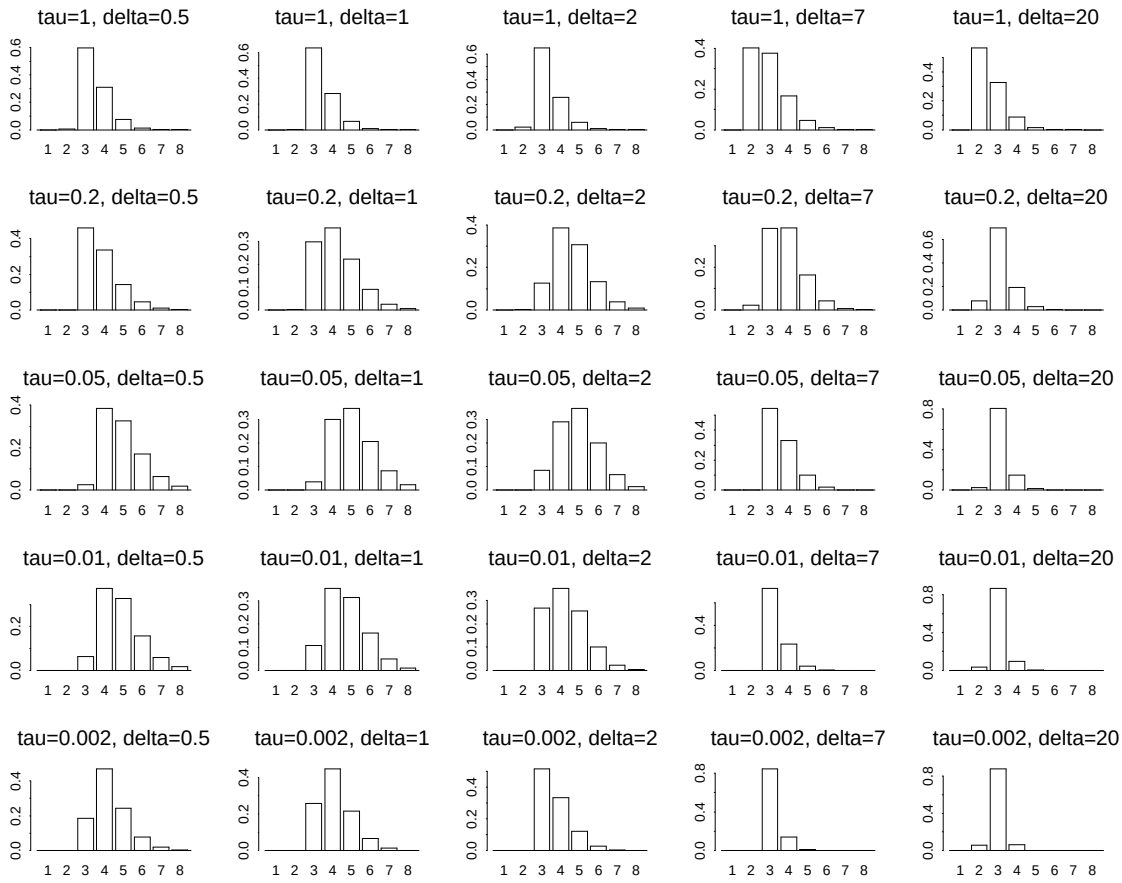


FIGURE 4: Posterior distribution of k for the galaxy data, using a $Poi(1)$ prior, for several values of the hyperparameters τ and δ.

The procedure I used to estimate τ and δ consists of taking the median of the posterior draws, after discarding those corresponding to values of k preceding the point where a rough level-off of the medians has occurred. The rationale is that if increasing k by 1 markedly changes the posteriors of τ and δ, it is because it affords a considerable reduction of the within-components variability, by replacing it with between-means variability. The median of τ seems to level off at $k = 4$. For δ the picture is less clear, but the decreases are much smaller past $k = 6$. The end result are the rough estimates $\hat{\tau} = 0.04$ and $\hat{\delta} = 2$. These values were used in the runs reported in Sections 2 and 3. A similar procedure was used by Nobile and Fearnside (2005), to which I refer for further examples.

The overall lesson seems to be that estimates of $\pi(k|x)$ provide only a rough, though useful, guide to the number of groups in the data and that there is really no substitute for the kind of sensitivity analysis performed in Figure 4.

Appendix

A.1 Proof of formula (14)

Letting $E = \mathcal{G}_t^\star$ in formula (12), with $t \leq k$, yields

$$f_k = \frac{\sum_{g \in \mathcal{G}_t^\star} f(g|k) f(x|k,g)}{\Pr[\mathcal{G}_t^\star | k, x]}.$$

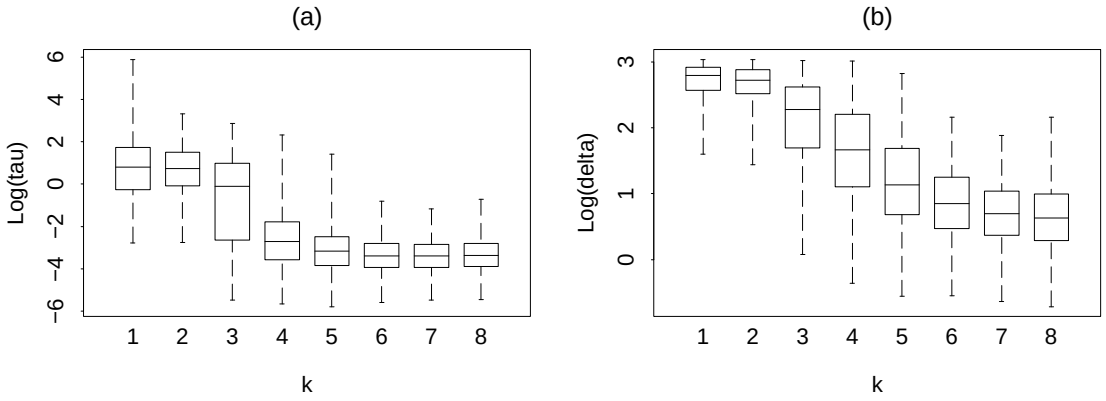


FIGURE 5: Boxplots of draws from the posterior distributions, conditional on k, of τ in panel (a) and of δ in panel (b). Whiskers are drawn at the 0.005 and 0.995 quantiles.

Now from formulae (5) the numerator equals $a_{kt} f_t^\star$, so that solving for $f_t^\star$ produces

$$f_t^\star = \frac{\Pr[\mathcal{G}_t^\star | k, x]}{a_{kt}} f_k.$$

Then taking the ratio between $f_{t+1}^\star$ and $f_t^\star$ and rearranging terms gives

$$f_{t+1}^\star \Pr[\mathcal{G}_t^\star | k, x] = a_{t+1,t}\, f_t^\star \Pr[\mathcal{G}_{t+1}^\star | k, x] \qquad t < k$$

where one uses $a_{k,t}/a_{k,t+1} = a_{t+1,t}$, a simple consequence of (5). Summing both sides over k from $t+1$ to $k_{\max}$ and rearranging yields (14).

A.2 Proof of formula (15)

Let $E = \widetilde{\mathcal{G}}_h^k$ in formula (12) to obtain

$$f_k = \frac{\sum_{g \in \widetilde{\mathcal{G}}_h^k} f(g|k) f(x|k,g)}{\Pr[\widetilde{\mathcal{G}}_h^k | k, x]} \qquad h \leq k.$$

The numerator is equal to $\binom{k}{h} a_{kh} f_h^\dagger$ (see Nobile 2004, Proof of Proposition 4.3), so that rearranging one has

$$f_h^\dagger = \frac{f_k \Pr[\widetilde{\mathcal{G}}_h^k | k, x]}{\binom{k}{h} a_{kh}}$$

Taking the ratio between $f_{h+1}^\dagger$ and $f_h^\dagger$ and rearranging terms produces

$$f_{h+1}^\dagger\, (k-h) \Pr[\widetilde{\mathcal{G}}_h^k | k, x] = f_h^\dagger\, (h+1)\, a_{h+1,h} \Pr[\widetilde{\mathcal{G}}_{h+1}^k | k, x].$$

Finally, sum both sides over $k = h+1, \ldots, k_{\max}$ and solve for $f_{h+1}^\dagger / f_h^\dagger$ to obtain (15).

A.3 Proof of formula (17)

From formula (5) and under the assumption that the prior is invariant with respect to permutations of the labels,

$$\binom{k}{t} a_{kt} = \binom{k}{t} \frac{\Gamma(k\alpha)}{\Gamma(k\alpha + n)} \frac{\Gamma(t\alpha + n)}{\Gamma(t\alpha)} \qquad t \leq k$$

Using formula 6.1.46 in Abramowitz and Stegun (1964), $n^{(k-t)\alpha}\dfrac{\Gamma(t\alpha+n)}{\Gamma(k\alpha+n)} \to 1$, as $n \to \infty$. Hence

$$n^{(k-t)\alpha}\binom{k}{t}a_{kt} \;\to\; \binom{k}{t}\frac{\Gamma(k\alpha)}{\Gamma(t\alpha)} \qquad \text{as } n \to \infty.$$

Therefore, for $t \le k$, $\binom{k}{t}a_{kt} \to \delta_{\{k=t\}}$ as $n \to \infty$.

A.4 Proof that every distribution satisfying equations (18) is truncated $Poi(1)$

The proof proceeds as follows: assume that $\pi(k) = 0$ for k larger than some value $\overline{k}$, use induction to derive $\pi(k)$ with $k \le \overline{k}$, finally let $\overline{k} \to \infty$. Suppose that, for all $k = j+1, \ldots, \overline{k}$, one has

$$\pi(k) = \pi(\overline{k})\frac{\overline{k}!}{k!} \tag{19}$$

Then equation (19) also holds for $k = j$:

$$\begin{aligned}
\pi(j) &= \sup_{k>j}\left\{\pi(k)\binom{k}{j}\right\} \\
&= \max_{j<k\le\overline{k}}\left\{\pi(\overline{k})\frac{\overline{k}!}{k!}\frac{k!}{j!(k-j)!}\right\} \\
&= \pi(\overline{k})\frac{\overline{k}!}{j!}
\end{aligned}$$

where the first line uses (18) while the second follows from (19) and $\pi(k) = 0$ for $k > \overline{k}$. Since equation (19) clearly holds for $k = \overline{k}$, an appeal to induction yields

$$\pi(k) = \pi(\overline{k})\frac{\overline{k}!}{k!} \propto \frac{1}{k!} \qquad k = 1, \ldots, \overline{k},$$

i.e., $Poi(1)$ restricted to $1 \le k \le \overline{k}$. Letting $\overline{k} \to \infty$ yields a $Poi(1)$ distribution restricted to $k \ge 1$.

References

Abramowitz, M. and Stegun, I. A. (1964). Handbook of Mathematical Functions, Dover edition, 9th printing, New York.

Aitkin, M. (2001). Likelihood and Bayesian analysis of mixtures. *Statistical Modelling*, **1**, 287–304.

Celeux, G., Hurn, M. and Robert, C. P. (2000). Computational and Inferential Difficulties with Mixture Posterior Distributions. *Journal of the American Statistical Association*, **95**, 957–970.

Chib, S. (1995). Marginal Likelihood from the Gibbs Output. *Journal of the American Statistical Association*, **90**, 1313–1321.

DiCiccio, T. J., Kass, R. E., Raftery, A. and Wasserman, L. (1997). Computing Bayes Factors By Combining Simulation and Asymptotic Approximations. *Journal of the American Statistical Association*, **92**, 903–915.

Diebolt, J. and Robert, C. P. (1994). Estimation of finite mixture distributions through Bayesian sampling. *Journal of the Royal Statistical Society B*, **56**, 363–375.

Frühwirth-Schnatter, S. (2001). Markov Chain Monte Carlo Estimation of Classical and Dynamic Switching and Mixture Models. *Journal of the American Statistical Association*, **96**, 194–209.

Gelman, A. and Meng, X.L. (1998). Simulating Normalizing Constants: From Importance Sampling to Bridge Sampling to Path Sampling. *Statistical Science*, **13**, 163–185.

Ishwaran, H., James, L. F. and Sun, J. (2001). Bayesian Model Selection in Finite Mixtures by Marginal Density Decompositions. *Journal of the American Statistical Association*, **96**, 1316–1332.

Nobile, A. (1994). Bayesian Analysis of Finite Mixture Distributions, Ph.D. dissertation, Department of Statistics, Carnegie Mellon Univ., Pittsburgh. Available at `http://www.stats.gla.ac.uk/~agostino`

Nobile, A. (2004). On the posterior distribution of the number of components in a finite mixture. *The Annals of Statistics*, **32**, 2044–2073.

Nobile, A. and Fearnside, A. (2005). Bayesian finite mixtures with an unknown number of components: the allocation sampler. Technical Report 05-4, Department of Statistics, University of Glasgow.

Phillips, D. B. and Smith, A. F. M. (1996). Bayesian model comparison via jump diffusions. In *Markov Chain Monte Carlo in Practice* (eds W. R. Gilks, S. Richardson and D. J. Spiegelhalter), 215–239, Chapman & Hall.

Raftery, A. E. (1996). Hypothesis testing and model selection. In *Markov Chain Monte Carlo in Practice* (eds W. R. Gilks, S. Richardson and D. J. Spiegelhalter), 163–187, Chapman & Hall.

Richardson, S. and Green P. J. (1997). On Bayesian analysis of mixtures with an unknown number of components (with discussion). *Journal of the Royal Statistical Society B*, **59**, 731–792.

Roeder, K. (1990). Density estimation with confidence sets exemplified by superclusters and voids in galaxies. *Journal of the American Statistical Association*, **85**, 617–624.

Stephens, M. (2000a). Bayesian analysis of mixture models with an unknown number of components – an alternative to reversible jump methods. *The Annals of Statistics*, **28**, 40–74.

Stephens, M. (2000b). Dealing with Label Switching in Mixture Models. *Journal of the Royal Statistical Society B*, **62**, 795–809.